On the Eyeshots of BGP Vantage Points

†Kai Chen, ‡Chengchen Hu, Wenwen Zhang, †Yan Chen, ‡Bin Liu
†Northwestern University, ‡Tsinghua University, University of Illinois at Chicago
†{kchen, ychen}@northwestern.edu, ‡{huc, liub}@tsinghua.edu.cn, wzhang31@uic.edu

***Abstract*—The publicly available BGP vantage points (VPs) have been heavily used by the research community to build the Internet autonomous system (AS) level topology, which is a key input to many applications, such as routing protocol design, performance evaluation and network security issues. However, a detailed study on the eyeshots[1] of these VPs has received little attention before. In this paper, we inspect these VPs carefully. Specifically, we do a measurement work to evaluate the effect of various factors on the eyeshot of each individual VP as well as the relationship between the eyeshots of different VPs. Based on the measurements, we disclose several counterintuitive observations and explain the possible reasons behind, which will help people to better understand the eyeshots of VPs and make better use of them in practice.**

***Index Terms*—Internet topology, Measurement, Vantage point, BGP routing table**

I. Introduction

The Internet AS topology is critical to many important applications. Nowadays, people's understanding about this topology is mainly based on several public route monitoring systems, such as Routeviews [1] and RIPE/RIS [2], which operate by gathering periodic BGP routing table snapshots and real-time BGP updates from ASes of various ISP backbones and network locations to discover dynamic changes of the global Internet routing system. Each of these peered ASes act as a VP for them to observe and construct the "entire" AS topology graph[2]. In normal peering contracts, these VPs are requested to regard the central route collectors as "customer" and configure their routers to export their full BGP routing table [4]. The central collectors are passive in the sense that they do not participate in announcing any prefixes or forwarding any traffic.

BGP-centric AS topological information collected and synthesized from all VPs has been heavily used by a substantial number of efforts [3, 5]–[15] to understand the Internet infrastructure, characterize its evolution, model its properties and infer its missing areas. Chang et al. [5] were among the first to study the completeness of commonly used BGP-derived topology maps. The pioneering efforts in [3, 6] were to evaluate and quantify the coverage of the public view on different components of Internet topology (*e.g.*, Tier-1 ASes, customer-providers links, and peering links) according to individual case studies. In [7], Oliveira et al. observed the tradeoff between topology liveness and completeness. In [8], the authors used BGP data to build a model to capture Internet routing diversity. He et al. [9] presented a framework to find the missing links in a snapshot of BGP-derived AS graph. In one seminal paper [13], Faloutsos et al. focused on the BGP-derived AS graph and argued that the AS topology follows several power-law rules, *e.g.*, its degree distribution. This observation was questioned in [14] by showing that the degree distribution, though heavy-tailed, did not follow strict power-law. Zhang et al. [15] examined the visibility constraints imposed by the deployment of route monitors on various applications. There were also efforts that attempt to describe the Internets AS-level topology with the help of generic graph-theoretic constructions (see for example [11, 12]).

Most previous work uses BGP data from a macroscopic view which simply blends the eyeshots of all the VPs together. To the best of our knowledge, the property of the eyeshot of each individual VP and the details of putting these eyeshots together to construct the AS graph have not been extensively studied before. In this paper, we focus on a microscopic view of studying the eyeshot of each individual VP. Intuitively, monitors placed in different VPs are supposed to provide different information about the global Internet structure. Questions such as what are the determinative factors of the eyeshot of an individual VP, what are the differences between the eyeshots of different VPs, and how the entire AS graph is assembled by all these eyeshots in a cumulative way, are interesting and are yet to be answered. To this end, we perform a comprehensive study on the eyeshots of these VPs. We make the following measurement efforts and contributions in this paper:

- We make a first detailed measurement study on the eyeshots of the public view VPs, by which we obtain counterintuitive and indicative observations.
- For the observations, we explain the reasons behind to help people to better understand them.
- Going beyond the measurements and observations, we discuss how our study could be used in practice to benefit the community.

The rest of this paper is organized as follows. Section II introduces the data collection and processing. Section III presents our measurements, observations and possible explanations. Section IV discusses the implications, applications as well as limitations. At last, Section V gives the concluding remark of the paper.

II. Data Collection and Processing

A. Dataset

To fully evaluate the eyeshot of a single VP, it is necessary for us to know all the backup links that can be observed by this VP. For this purpose, the BGP data we used in this paper include a collection of BGP routing tables/updates from 790 BGP speaking routers in 438 unique ASes from Dec, 2007 to Sep, 2008. Specifically, we start from the dataset of Routeviews [1] collected at route-views.oregon-ix.net, which

[1]By the "eyeshot" of a VP, we mean the AS links this VP has observed and reported to the central collectors.

[2]It is known that the BGP-based AS topology is not complete, we call it entire because it is a best effort to date [3].

is the most widely used BGP archive so far. We then include routing data from 6 other Oregon route severs and 16 route collectors of RIPE/RIS [2]. We term this dataset as "public view". According to previous researches [3, 6, 7], the ten-month public view data are enough to cover all the backup links in the Internet topology. Though there are additional BGP data sources such as route severs, looking glasses, and Internet Routing Register [16], the amount of additional AS links they could add upon to the public view is quite incremental [3, 6]. Additionally, some of these additional datasets do not provide historical archives which further limit their information on the backup links. So we do not include them here.

B. Policy Inference

After collecting and extracting AS paths as described previously, we leverage the previous work [17] to infer the relationship between ASes. Although there are several follow-up works on relevant policy inference [18]–[20], we prefer [17] due to its ability of inferring relationships based only on the AS paths, without any other extra requirements such as data from Internet registers or active probes. While simple, the author showed the 99.1% correctness of their inferred relationships between AT&T and its neighbors. There are also comparison studies of the accuracy of algorithms in [17] and [19] indicating that the former one is more accurate in identifying peer-peer relationship. Furthermore, study in [15] has shown that the AS relationships inferred by heuristics in [17] are pretty stable with respect to variations in the observed AS paths. In addition, we assume that the AS relationships did not change significantly within a ten-month period. To justify this, we sample the AS relationships from CAIDA [21] for the past five years. We check the relationships at ten-month intervals and find that more than 98.5% of AS pairs do not change their relationships. The statistics of the inference result is listed in Table I.

# Total	# Customer-provider	# Peer-peer	# Others
117580	90755 (77.2%)	18687 (15.9%)	8138 (6.9%)

TABLE I
STATISTICS OF LINKS OF THE PUBLIC VIEW (THE "OTHERS" INCLUDES THE SIBLING-SIBLING LINKS AND LINKS WHOSE RELATIONSHIPS HAVE NOT BEEN DETERMINED).

C. Network Classification

There are several ways to classify the tiers of ASes. Existing approaches use the degree of an AS, the number of prefixes originated by an AS, or the distinct AS paths seen from an AS, etc. However, without considering the AS relationships, these heuristics may be misleading [6]. For instance, the degree for an AS could be a mixed set of neighbors including providers, peers, or customers. Algorithms using the number of prefixes and distinct AS paths may be too coarse since prefixes vary significantly in length and route aggregations happen everywhere. To mitigate these limitations, we use the method proposed in [6]. This method uses the number of downstream customer ASes to classify AS hierarchy. According to this rule, we classify all the ASes into five tiers as shown in Table II.

# Tier-1	# Tier-2	# Tier-3	# Tier-4	# Tier-5
9	479	1588	3454	26233

TABLE II
STATISTICS OF ASES ON EACH TIER.

III. MEASUREMENTS AND OBSERVATIONS

In this section, we perform a detailed measurement study on the eyeshots of all VPs. First, we measure the eyeshot of each individual VP to understand the relation between its eyeshot and various factors. Then, we focus on characterizing eyeshots of every pair of VPs and investigating if the AS relationship between two VPs has impact on the differences between their eyeshots. Finally, we examine how the entire AS graph is cumulatively constructed by the eyeshots of VPs. We also explain the possible reasons behind some of the observed phenomena.

A. Individual VP Eyeshot

We check the effects of several factors on the eyeshot of each individual VP and outline the determinative factor for estimating the eyeshot of a single VP.

Total #	Tier-1	Tier-2	Tier-3	Tier-4	Tier-5
438	9	116	112	110	91

TABLE III
STATISTICS OF VPS ON EACH TIER.

Rank	AS number	AS name	# of AS links	Tier
1	11608	ACTTG	68958	2
2	3277	RUSNET	66353	4
3	12989	HWNG	64333	2
4	1280	ISC	63926	2
5	4513	INTERNAP-4513	63350	2
6	286	KPN	63157	2
7	14361	HOPONE-global	63114	3
8	7575	AARNET	62941	2
9	29073	ECATEL	62743	5
10	2497	IIJ	62660	2

TABLE IV
THE TOP 10 VPS AND THEIR EYESHOTS.

1) Relation with Tier: An intuitive thought that the eyeshot of a particular VP could have relation with the tier where it stands and the higher the tier, the better its eyeshot, is actually not true. To evaluate this, we classify all the VPs according to their network tiers (as shown in Table III) and check if the eyeshot has correlation with the tier the VP lies in. Fig. 1 gives the individual value plot of the eyeshots for all tiered VPs. From the figure, we can clearly see that while all Tier-1 VPs can observe a substantial number of links in general, VPs on all the other tiers have a high variability on their eyeshots. For example, there exist Tier-2 VPs who can observe very few links; on the contrary, there exist Tier-5 VPs who can capture more than 60,000 links. We also annotate the mean and median values in the figure, and find that both decrease with tiers. It should be noted that the median values for Tier-3, Tier-4, and Tier-5 VPs approach 0 which means that most of the VPs in these tiers only observe little links, say one or two links. Furthermore, Table IV enumerates the top ten VPs according

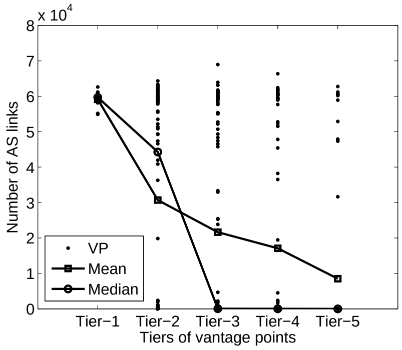

Fig. 1. Eyeshot vs Tier.

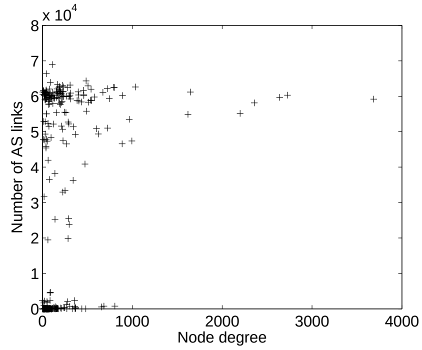

Fig. 2. Eyeshot vs Degree.

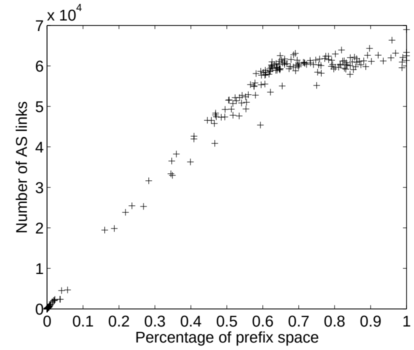

Fig. 3. Eyeshot vs Prefix.

to their eyeshots. It is surprising that there is no Tier-1 VP within these ten ASes. While most of them are Tier-2 VPs, we witness Tier-4 and Tier-5 VPs among this list.

- ***Observation 1:*** The mean and median values of the eyeshots decrease with tiers, however the maximum values have no relationship with tiers, in other words, tier cannot be a determinate measure to show the eyeshot of a VP.

2) Relation with AS degree: A VP will learn the routes to each destination from all its neighboring ASes. So intuitively, VPs with more neighbors may be able to see more links. To this end, we study if the AS degree of a VP can be used to predict the eyeshot of this VP. Fig. 2 demonstrates the relation between the degree of a VP and the number of AS links it can observe. We find a poor correlation from this figure. While a high degree (when the AS degree is greater than 1,000) usually indicates a good eyeshot, there is a high variability among these middle degree and low degree VPs. For instance, while there exists some VPs of very low degree (tens) who can provide a substantial number of links, there are some other VPs of middle degree (hundreds) but only contribute very limited number of links to the public view.

- ***Observation 2:*** The eyeshot of a VP with a large AS degree is always large; however, the eyeshot of a VP with a small AS degree is not necessarily small.

3) Relation with the observed prefixes: According to the global reachability property of the Internet routing, each VP should be able to receive the route(s) to every routable prefix. However, due to issues such as route aggregation and policy filtering, one VP may only observe a restricted number of prefixes. We collect all the prefixes that are observed by each VP and draw a scatterplot between the observed prefixes (normalized by the whole prefix spaces) and the eyeshot of this VP in Fig 3. We find that the number of AS links each VP can contribute is highly correlated with the number of observed prefixes especially when the percentage is less than 65%. This correlation implies that the number of observed prefixes is a good indicator to predict the eyeshot of a VP. Further, when the percentage is larger than 65% the eyeshots remains stable.

- ***Observation 3:*** The number of observed prefixes is one determinative measure for the eyeshot of a particular VP. Up to 65%, the eyeshot of a VP is highly (linearly) correlated with the number of prefixes it has observed. However, when exceeding this value, the eyeshot remains stable around 60,000 no matter how many prefixes it can observe.

4) Explanation: We speculate there are following reasons accounting for the above results, *i.e.*, why the tier or AS degree of a VP cannot be a good measure for its eyeshot? First, the existence of route aggregations across the Internet heavily narrows the eyeshots of some VPs. And the random occurrence of aggregations may partially explain the high variability of the eyeshots of VPs on the same tiers. Second, some VPs do not provide the full table as they are supposed to. This renders it difficult to see some links from these VPs. For instance, a VP may regard the central collector as "peer" instead of "customer", and will only export the routes from its customers but not the full table. To validate this conjecture, we directly contact the Routeview administrator and get positive confirmation [4]. Third, it is possible that some VPs have not been able to capture all the backup links. However, this may not be a major reason given that ten-month public view data have a very good coverage on the backup links [7].

B. Comparison between VPs

Located in different places, each VP should have its own view about the Internet topology. Therefore, we compare the eyeshots of any two VPs, quantify the differences and quest the reasons behind the observations.

1) Selecting representative VPs: As stated above, ideally, each VP should be able to see all known address spaces for the global reachability. However, due to issues such as prefix aggregation and filtering, the complete address spaces cannot be seen by all VPs. We differentiate the VPs according to the following three measures: 1) the number of observed prefixes, 2) the number of observed ASes, and 3) the number of the observed links. When ordering the VPs in terms of the number of observed prefixes, we find all these three measures dramatically decrease between the 100th VP and the 200th VP, especially around the 150th VP. So we conservatively choose the first 100 VPs for the analysis in this section because these VPs have a good coverage of address spaces, in other words, they are representative.

- ***Observation 4:*** Although there are more than 400 VPs in the public view, most of them contribute very little.

2) Quantifying the differences: To quantify the dissimilarity between two VP's eyeshots, we measure the *Jaccard distance* between them. As shown in (1), this is calculated by dividing the sizes of the differences of two sets by the size of the union. Note that $0 \leq J \leq 1$ measures how different these two eyeshots

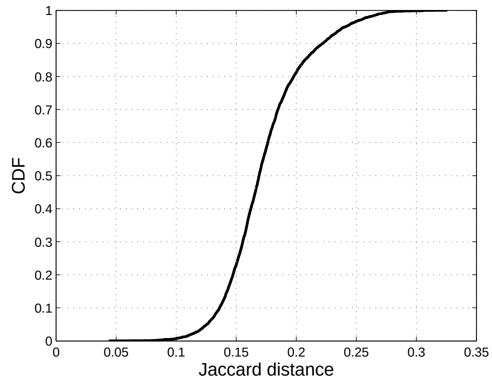

Fig. 4. CDF plot of the *Jaccard distance* among all VP pairs.

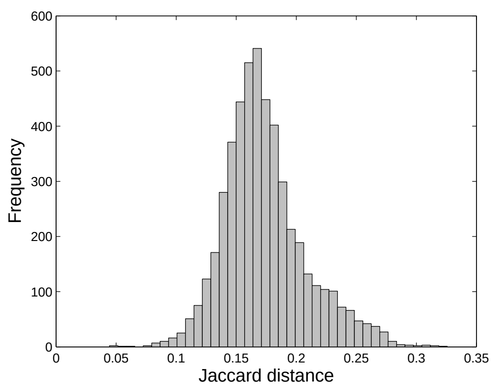

Fig. 5. Histogram plot of *Jaccard distance* among all VP pairs.

are. For example, a distance of 0.1 indicates that 10% of the links are seen from either VP_1 or VP_2, but not both.

$$J(s_1, s_2) = \frac{|(s_1 - s_2) \cup (s_2 - s_1)|}{|s_1 \cup s_2|} \tag{1}$$

Focusing on the 100 selected representative VPs, in Fig. 4, we show the cumulative distribution function (CDF) of *Jaccard distances* of these $\binom{100}{2} = 4950$ VP pairs. From the CDF, it is interesting to find that the curve is approximately normal with a central tendency $\mu = 0.174$ and standard deviation $\sigma = 0.034$. To further display this property, Fig. 5 shows the histogram which can give a more clear sense about the approximate normality. Through this, we can also learn that, on an average, the eyeshots of any two representative VPs can have as high as 82.6% links in common.

- ***Observation 5:*** The *Jaccard distances* of the eyeshots of representative VP pairs approximately follow a nice normal distribution, *i.e.*, $N(0.174, 0.034^2)$.

3) Impact of peering relationship: We now study how the peering relationship affects the *Jaccard distances* between the VP pairs. Among these 4950 pairs, 3960 pairs do not connect with each other, 457 pairs have *peer-peer* relationships, and 533 pairs have *customer-provider* relationships. As we check, the average *Jaccard distance* of unconnected pairs is 17.8%, while the average value of *peer-peer* pairs is 15.6% and of *customer-provider* pairs is 15.7%. It makes sense that VP pairs that do not have connections cannot exchange routes directly, so averagely, they tend to have less in common between their eyeshots than those having connections. In addition, we find the average *Jaccard distance* of *peer-peer* pairs and that of *customer-provider* pairs are too close to be differentiated.

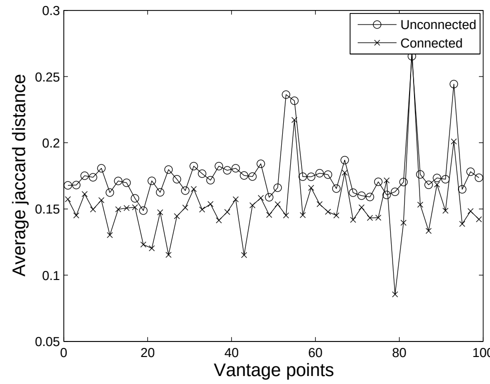

Fig. 6. Average *Jaccard distances* between connected and unconnected VPs.

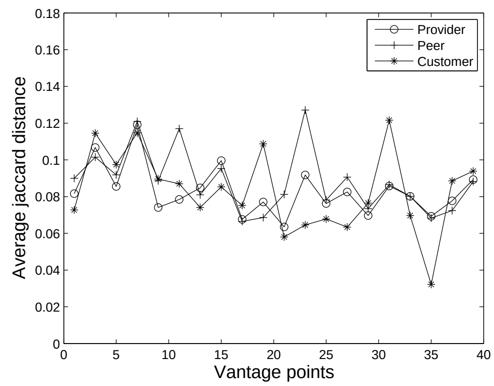

Fig. 7. Average *Jaccard distances* with providers, peers, and customers.

- ***Observation 6:*** The eyeshots of VP pairs who are not connected have less in common than those who are connected.

As a further step, we study from the perspective of each individual VP. We calculate and plot the average *Jaccard distance* of each VP with all its unconnected VPs and that with all its connected ones in Fig. 6. The results confirm the previous observation that a VP tends to have more in common with its neighboring VPs than those unconnected ones. We further measure the average *Jaccard distances* of each VP with its *providers*, *peers*, and *customers* in Fig. 7. Note that there are 39 VPs having *providers*, *peers* and *customers* simultaneously, so we only include them in the figure. We do not see any obvious difference among these three categories. According to the valley-free routing policy [17], an AS will export most of its known to its *customers*, less to its *peers* or *providers*, so it is expected that the eyeshot of one VP may be covered by its neighbors differently. However, our experiments show that this is not the case. That is to say, the eyeshot of a VP is covered by its neighbors in a relatively random way regardless of their mutual AS relationship. We believe route aggregation is one root cause for this phenomenon. Ideally, a provider should export all its routes to its customer, so the *Jaccard distance* should be zero. However, due to aggregations, the provider may have aggregated some routes before exporting to the customer, which result in the difference between their eyeshots.

- ***Observation 7:*** The eyeshot of a VP is covered by its neighbors in a relatively random way regardless of their mutual AS relationship.

C. Assembling the Entire AS Graph

We check how the AS graph is cumulatively constructed by the eyeshots of VPs. In addition to the first 100 VPs, we add another 50 VPs according to Section III-B1. Though these VPs do not provide a complete coverage of address spaces, the still contribute to our study because they have observed amount of links. Note that we choose VPs using the following orders:

- *Link-based selection*: VPs are chosen based on the eyeshots, and VPs with more observed links are preferred to those with fewer ones.
- *Random-based selection*: VPs are chosen randomly.

1) Marginal coverage: Before directly jumping into assembling the AS graph, we are interested in the question: how much link information of a given VP can also be observed by the other VPs? To answer this question, for each VP we calibrate how many other VPs can cover the eyeshot of the chosen one. Our experimental result exhibits that most of the eyeshots of VPs can be well covered by tens of other VPs, and some of them only need less than ten other VPs. For illustration, in Fig. 8, we take a well-known Tier-1 VP, AT&T (AS7018), as an example to demonstrate the coverage process, and other VPs have similar curves with this one (We do not include them here for the space limitation). We can clearly see that this VP's eyeshot can be fully covered by 20 others VPs even in a random order.

- ***Observation 8:*** Most of the eyeshots of VPs can be well covered by tens of other VPs, and some of them are even covered by less than ten other VPs.

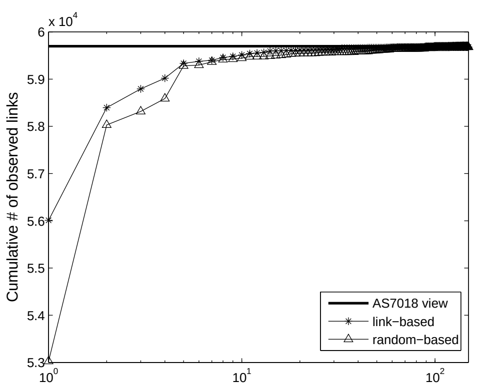

Fig. 8. Marginal utility on Tier-1 AS7018.

2) Convergence process: In Fig. 9, we examine the convergence process by incrementally adding eyeshots of VPs. A key observation here is that the random scheme is comparable with the link-based scheme in building this AS topology graph. Approximately, 50 VPs can assemble 80% of the entire graph, and 100 VPs can assemble more than 90%.

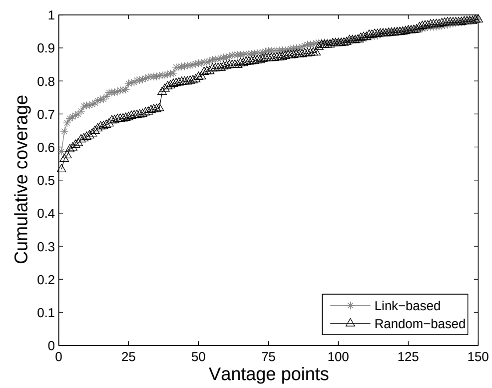

Fig. 9. Convergence process by cumulatively adding VPs.

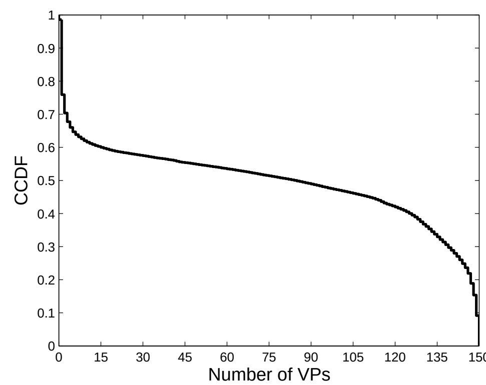

Fig. 10. Links VS the number of VPs observe them simultaneously.

- ***Observation 9:*** Different ways of selecting VPs do not have many differences on the assembly convergence. Furthermore, randomly choosing 100 VPs can assemble more than 90% of the entire AS topology graph.

3) Explanation: To explain why the random scheme has good performance as the link-based one, we calculate for each AS link how many VPs among these 150 selected ones have seen it simultaneously. In Fig. 10, we draw the complementary cumulative distribution function (CCDF). It is obvious that most of the links can be observed from multiple VPs. For example, more than 50% of the total links can be seen from the eyeshots of 75 VPs. To its extreme, there are around 10% of the links can be seen from the eyeshots of all 150 VPs. This explains why the random scheme could have such good performance. Since one link can be captured from multiple VPs, one still has good chance to observe it when randomly choosing a VP.

IV. Discussion

Although the data collected from all the public view VPs have been heavily used to form the basis for many important researches, the eyeshot of each individual VP has not been intensively studied before. We posit that a thorough investigation on the eyeshots of VPs will help people to better understand them and make better use of them in practice. Specifically, the measurements and observations in Section III could have the following implications and applications.

First, the operation and maintenance of the VPs for the measurement infrastructure require lot of money each year. Several measurement projects or infrastructures finally end up with the shortage of support funding. For example, NLANR is a concrete lesson to learn [22]. From our **observations 8** and **9**, the eyeshots of the current VPs have large redundancy. It is possible to reduce the number of the VPs while provide similar measurement power. In this way, the annual cost for the VPs can be greatly saved. Therefore, the study of how to reduce the redundancy of VPs while provide the required measurement power should be an attractive and practical topic.

Second, our characterization on the eyeshots from different VPs, especially the identified normal distribution (*i.e.*, **observations 5**), could be an useful input to construct a mathematical model to infer how many missing links are there in the ground-truth of AS topology.

Third, from the perspective of common users, our measurement results, such as **observations 4, 8** and **9**, can be referred to save their labours by collecting the data from less VPs without losing much topological information. Of course, it is required to pick out a minimum set of VPs that can provide users such a shortcut in their efforts. From the perspective of management of the public view, our study is not only a good starting point for designing an efficient deployment of VPs, but also a good estimator to estimate the potential of a newly coming VP (such as lessons learned from **observations 1, 2,** and **3**).

Fourth, the greening of Internet has drawn increasing attention in recent years. It is a general trend that the Internet should operate more greenly in future. Our observations in the above section serve as a good feedback to the relevant administrators suggesting that certain actions could be taken to make their monitoring infrastructures more energy-efficient and green.

Though having positive indications, our study also has its own limitations. For instance, we are focusing on the static information, *i.e.*, AS links, in this paper. Although AS link is a most important piece of information in the AS topology study, there is other information such as dynamics that also need to be attended in our future work. Furthermore, the experimental results are customized to our dataset from Dec, 2007 to Sep, 2008. While the results are expected to be held at a different time window, a detailed evaluation should be carried out before we make our results widely applicable. All these limitations will be considered in our coming work.

V. Conclusion

This paper differentiates itself from all previous work by a first step to detailedly study the eyeshots of the public view VPs. First, we found that the eyeshot of an individual VP is poorly correlated with its tier or degree which may be under intuitive assumptions, instead, the number of observed prefixes is one indicator for the eyeshot. Second, we observed a nice normal distribution of the *Jaccard distances* of the eyeshots between the representative VPs. We also identified that the mutual AS relationship has no obvious effect on the dissimilarity of the eyeshots. Third, we found the two different schemes of selecting VPs do not have much differences on the process of building the AS topology graph. We further showed that the eyeshots of these VPs have great redundancy so that randomly choosing 100 of them can assemble more than 90% of the entire AS topology graph. For some observations we made, we have specifically explained the possible reasons behind for people's better understanding. At last, we have discussed the implications, applications as well as limitations of this study, and the discussion actually formulated our forthcoming agenda.

Acknowledgment

This work is supported by grants from NU-Motorola Center for Seamless Communication, NSF awards CT-0627715 and CNS-0613967, NSFC (60625201, 60873250), the Specialized Research Fund for the Doctoral Program of Higher Education of China (20060003058), 863 high-tech project (2007AA01Z216, 2007AA01Z468) and open project of state key Laboratory of Networking and Switching Technology (SKLNST-2008-1-05).

References

[1] University of Oregon Route Views Archive Project, "http://www.routeviews.org/," .

[2] RIPE, "Routing Information Service," http://www.ripe.net/projects/ris/.

[3] R. Oliveira, D. Pei, W. Willinger, B.Zhang, and L. Zhang, "In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure," in *ACM SIGMETRICS*, 2008.

[4] Authors and Administrator, "Question about Routeviews," in *Email communication between administrator of Routeview project*, 2008.

[5] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," in *Computer Netw.*, 2004.

[6] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE Tran. on Networking*, 2010.

[7] R. Oliveira, B. Zhang, and L. Zhang, "Observing the Evolution of Internet AS Topology," in *ACM SIGCOMM*, 2007.

[8] W. Mhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-Topology Model that Captures Route Diversity," in *ACM SIGCOMM*, 2006.

[9] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A Systematic Framework for Unearthing the Missing Links: Measurements and Impact," in *NSDI*, 2007.

[10] P. Mahadevan, D. Krioukov, M. Fomenkov, and B. Huffaker, "The Internet AS-level topology: Three Data sources and One definitive metric," in *ACM SIGCOMM CCR*, 2006.

[11] Priya Mahadevan, Calvin Hubble, Dmitri Krioukov, Bradley Huffaker, and Amin Vahdat, "Systematic Topology Analysis and Generation Using Degree Coorelations," in *ACM SIGCOMM*, 2006.

[12] Priya Mahadevan, Calvin Hubble, Dmitri Krioukov, Bradley Huffaker, and Amin Vahdat, "Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies," in *ACM SIGCOMM*, 2007.

[13] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-law Relationships of the Internet Topology," in *ACM SIGCOMM*, 1999.

[14] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The Origin of Power Laws in Internet Topologies Revisited," in *IEEE Infocom*, 2002.

[15] Y. Zhang, Z. Zhang, Z. Morley Mao, Y. Charlie Hu, and B. Maggs, "On the impact of route monitor selection," in *Proceedings of ACM Internet Measurement Conference*, 2007.

[16] "Internet Routing Register," http://www.irr.net.

[17] Lixin Gao, "On inferring Autonomous System relationships in the Internet," *IEEE/ACM Trans. Netw.*, 2001.

[18] Jianhong Xia and Lixin Gao, "On the Evaluation of AS Relationship Inferences," in *In IEEE GLOBECOM*, 2004.

[19] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet Hierarchy from multiple Vantage Points," in *Proc. of IEEE Infocom*, 2002.

[20] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K.C. Claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM CCR*, 2007.

[21] "The CAIDA AS Relationships Dataset, 2004-2008," http://www.caida.org/data/active/as-relationships/.

[22] "National Laboratory for Applied Network Research," http://www.nlanr.net/.